FROM TEXT TO TALK: AUDIO-LANGUAGE MODEL NEEDS NON-AUTOREGRESSIVE JOINT TRAINING

Anonymous authors

000

001

002003004

010 011

012

013

014

015

016

017

018

019

021

025

026

027 028 029

030

032

033

034

037

040

041

042

043

044

046

047

048

051

052

Paper under double-blind review

ABSTRACT

Recent advances in large language models (LLMs) have attracted significant interest in extending their capabilities to multimodal scenarios, particularly for speechto-speech conversational systems. However, existing multimodal models handling interleaved audio and text rely on autoregressive (AR) methods, overlooking that text depends on target-target relations whereas audio depends mainly on sourcetarget relations. In this work, we propose Text-to-Talk (TtT), a unified audio-text framework that integrates AR text generation with non-autoregressive (NAR) audio diffusion in a single Transformer. By leveraging the any-order AR property of absorbing discrete diffusion, our approach provides a unified training objective for text and audio. To support this hybrid generation paradigm, we design a modality-aware attention mechanism that enforces causal decoding for text while allowing bidirectional modeling within audio spans, and further introduce three training strategies that reduce train-test discrepancies. During inference, TtT employs block-wise diffusion to synthesize audio in parallel while flexibly handling variable-length outputs. Extensive experiments across Audio-QA and ASR tasks demonstrate the effectiveness of our approach, with detailed ablation studies validating each proposed component. We will open-source our models, data and code to facilitate future research in this direction.

1 Introduction

The recent success of LLMs has catalyzed a paradigm shift towards general-purpose Multimodal Large Language Models (MLLMs) capable of processing and generating information across diverse modalities (Xu et al., 2025; Team et al., 2023). Among these, speech-to-speech conversational systems have emerged as a pivotal component in facilitating natural human-AI interaction. Conventional systems typically decompose this problem into a cascaded pipeline of Automatic Speech Recognition (ASR), LLM-driven response generation, and Text-To-Speech (TTS) synthesis. While effective to a degree, this modular design introduces significant latency accumulation and error propagation between modules, hindering naturalness and real-world applicability. In response, recent end-to-end approaches like Moshi (Défossez et al., 2024), GLM4-Voice (Zeng et al., 2024), and VITA-Audio (Long et al., 2025) have sought to unify speech understanding and generation within a single model. These models are typically trained through multi-stage pipelines that involve textto-audio tokenizer training, interleaved data construction, text-audio alignment and task-oriented supervised fine-tuning (Huang et al., 2025; Li et al., 2025; Ding et al., 2025; Chu et al., 2024). As shown in Figure 1, these methods aim to generate interleaved text and speech tokens in an autoregressive (AR) manner, which are then decoded into continuous audio waveforms by a separate neural codec or diffusion-based decoder (Mehta et al., 2024; Kong et al., 2020).

However, this emerging paradigm faces a fundamental challenge. As illustrated in Figure 1, we identify a fundamental mismatch in prevailing approaches that employ a single language model to autoregressively generate both text and audio tokens (Zeng et al., 2024; Xie & Wu, 2024; Borsos et al., 2023; Dang et al., 2024; Rubenstein et al., 2023). This uniform treatment applies identical AR training objectives across both modalities, overlooks a critical distinction in their underlying generative processes. Text generation inherently follows a sequential causal structure characterized by strong **target-target** dependencies (Box et al., 2015), where each token explicitly conditions on previously generated tokens. Consequently, an incorrect token prediction can propagate and introduce subsequent errors due to the exposure bias inherent in AR models (Ranzato et al., 2015). In

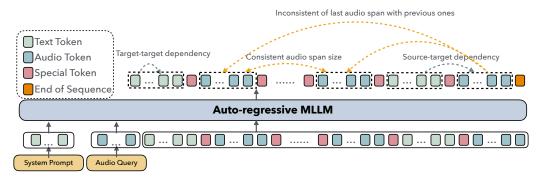


Figure 1: (a) Distinct dependency structures for text and audio modality. (b) Due to disparate tokenization rates, the last audio span is of variable length.

contrast, audio token generation is predominantly driven by **source-target** dependencies (Ren et al., 2020), where audio output primarily condition on the source text rather than on the preceding audio tokens. Specifically, within the current non-autoregressive (NAR) span, audio tokens generation should remain faithful to the source text even when previous audio tokens are incorrectly predicted. Applying a purely AR objective to audio generation thus introduces unnecessary sequential constraints, leading to suboptimal training dynamics and magnifying error propagation. This problem can be substantially alleviated by adopting a NAR generation strategy, which aligns better with the source-dependent nature of audio modeling. Recently, discrete diffusion has emerged as a compelling alternative to AR for discrete sequence modeling (Yu et al., 2025; Gong et al., 2024; Austin et al., 2021; Sahoo et al., 2024). Beyond empirical gains, recent theory shows that absorbing discrete diffusion can be interpreted as modeling the conditional distributions of clean tokens and admits a tight connection to any-order AR objectives (Ou et al., 2024).

Thus, we introduce Text-to-Talk (TtT), a unified audio-text MLLM that integrates AR text generation with NAR audio diffusion within a single Transformer initialized from a pretrained LLM. Text segments are trained with a standard AR cross-entropy objective, while audio segments are modeled via an NAR discrete diffusion process. During inference, the model dynamically switches between AR and NAR decoding strategies based on special control tokens. In summary, our work makes the following contributions:

- We identify and formalize the fundamental asymmetry in dependency structures between text and audio modalities—text exhibits target-target dependencies requiring causal ordering, while audio is driven by source-target dependencies. Leveraging the any-order AR nature of absorbing discrete diffusion, we establish a unified theoretical framework that proves our joint training objective provides an upper bound on the negative log-likelihood of the desired joint distribution.
- We propose TtT, a hybrid AR-NAR MLLM that seamlessly integrates AR text generation with discrete diffusion-based audio synthesis within a single Transformer initialized from a pretrained LLM. Our design preserves the reasoning and instruction-following capabilities of the base LLM while enabling efficient parallel audio generation.
- We introduce three principled training strategies to address the inherent train-test discrepancies in hybrid AR-NAR learning, enabling stable training and robust content-aware variable-length generation that bridges the gap between training and inference conditions.
- Extensive experiments on Audio Question Answering (Audio-QA) and ASR tasks show that our proposed TtT framework substantially outperforms strong AR and NAR baselines. Our results underscore the superiority of the hybrid AR-NAR architecture, and detailed ablations validate the effectiveness of our specialized training strategies.

2 Preliminary and Notation

In this section, we establish the basic notation for interleaved audio-text sequences and provide brief overviews of the two core generative paradigms employed in our framework: AR modeling and absorbing discrete diffusion. These form the theoretical foundation of our proposed method in Section 3.

Tokens, Vocabulary, and Interleaved layout We consider interleaved discrete text–audio sequences of length L: $x = (x^1, \dots, x^L)$ with a unified discrete vocabulary $\mathcal{V} = \mathcal{V}_{\text{text}} \cup \mathcal{V}_{\text{audio}} \cup \mathcal{S}$, where \mathcal{S} contains special tokens such as $\langle \text{SOA} \rangle$ (start of audio), $\langle \text{EOA} \rangle$ (end of audio), $\langle \text{EOS} \rangle$ (end of sequence) and the absorbing mask token $[\mathbf{M}]$. A sequence x is structured as a series of alternating text and audio spans: $x = (\mathcal{T}_1, \mathcal{A}_1, \dots, \mathcal{T}_M, \mathcal{A}_M, \langle \text{EOS} \rangle)$, where:

- $\mathcal{T}_m = (t_{m,1}, \dots, t_{m,|\mathcal{T}_m|}) \in (\mathcal{V}_{\text{text}} \cup \{\langle \text{EOS} \rangle, \langle \text{SOA} \rangle\})^{|\mathcal{T}_m|}$ are text tokens
- $\mathcal{A}_m = (a_{m,1}, \dots, a_{m,|\mathcal{A}_m|}) \in (\mathcal{V}_{\mathrm{audio}} \cup \{\langle \mathrm{EOA} \rangle\})^{|\mathcal{A}_m|}$ are quantized audio tokens

Let $f_{\theta}: \mathcal{V}^{L} \to \mathbb{R}^{L \times d}$ be a single Transformer (e.g. Qwen 2.5). We use a shared output head $W \in \mathbb{R}^{d \times |\mathcal{V}|}$ (typically tied with input embeddings) to produce per-position logits over the entire vocabulary \mathcal{V} .

AR Modeling AR models factorize the joint probability of a sequence $x = (x^1, \dots, x^L)$ into a product of conditional probabilities, based on the chain rule: $p(x) = \prod_{i=1}^L p(x^i|x^{< i})$, where $x^{< i} = (x^1, \dots, x^{i-1})$. This imposes a sequential, causal structure on the generation process. For a detailed discussion, please refer to Appendix A.3.1.

Absorbing Discrete Diffusion Absorbing discrete diffusion models are a NAR paradigm for sequence generation. They consist of a forward process that corrupts a clean sequence by gradually replacing tokens with a special absorbing mask state [M], and a learned reverse process that aims to recover the original sequence from the corrupted input. A key insight from Ou et al. (2024) is that the learning objective simplifies to modeling a time-independent conditional probability of the clean data. Specifically, the score for unmasking a token v at a corrupted position is given by:

$$\underbrace{\frac{p_t(\dots,\hat{x}^i = v,\dots)}{p_t(\dots,x^i = [\mathbf{M}],\dots)}}_{\text{concrete score}} = \underbrace{\frac{e^{-\bar{\sigma}(t)}}{1 - e^{-\bar{\sigma}(t)}}}_{\text{time scalar}} \cdot \underbrace{p_0(v \mid UM)}_{\text{clean conditional probability}}.$$
 (1)

where UM denotes the set of unmasked (visible) tokens and t represents the continuous time step of the corruption process. This denoising formulation is precisely the objective of an Any-Order AR Model (AO-ARM), predicting a token given an arbitrary context of unmasked tokens. As demonstrated by Ou et al. (2024), the diffusion training objective is mathematically equivalent to this AO-ARM objective, which averages the prediction loss over all possible permutations of the sequence: $\mathcal{L}_{AO}(\boldsymbol{x}_0) = \mathbb{E}_{\pi \sim U_\pi} \sum_{l=1}^d -\log q_\theta(x_0^{\pi(l)}|x_0^{\pi(<l)})$, where π is a random permutation of the token indices. Therefore, training an absorbing discrete diffusion model is equivalent to training a powerful ensemble of AR models that can operate in any order. More details in Appendix A.3.2.

3 JOINT TEXT-AR & AUDIO-NAR MODEL

In this section, we introduce our proposed model, integrates AR generation for text and discrete diffusion for audio within a single, unified Transformer architecture.

3.1 AR Modeling for Text

We model text generation using a fixed, canonical auto-regressive order. Let π_{text} denote the natural left-to-right permutation over all text token positions in the sequence — that is, $\pi_{\text{text}}(1) < \pi_{\text{text}}(2) < \cdots < \pi_{\text{text}}(|\mathcal{T}_{\leq M}|)$ where $\mathcal{T}_{\leq M} = \cup_{m=1}^{M} \mathcal{T}_m$ is the set of all text token indices.

At the span level, the probability of generating the m-th text span $\mathcal{T}_m = (t_{m,1}, \dots, t_{m,|\mathcal{T}_m|})$ conditioned on all prior context is given by: $p_{\theta}(\mathcal{T}_m | \mathcal{T}_{< m}, \mathcal{A}_{< m}) = \prod_{j=1}^{|T_m|} p_{\theta}(t_{m,j} | \mathcal{T}_{< m}, \mathcal{A}_{< m}, t_{m,< j})$, where $t_{m,< j} = (t_{m,1}, \dots, t_{m,j-1})$ is the prefix of text tokens within the current text span.

To express the joint probability of all text tokens in the sequence, we account for the conditioning on preceding audio spans. The joint probability $p_{\theta}(x_{text})$ is therefore defined as the product of the probabilities of each text span, conditioned on all prior spans: $p_{\theta}(x_{text}) = \prod_{m=1}^{M} p_{\theta}(\mathcal{T}_m | \mathcal{T}_{< m}, \mathcal{A}_{< m}) = \prod_{m=1}^{M} \prod_{j=1}^{|\mathcal{T}_m|} p_{\theta}(t_{m,j} | \mathcal{T}_{< m}, \mathcal{A}_{< m}, t_{m,< j}),$

The model is trained by minimizing the standard causal cross-entropy loss over all text positions:

$$\mathcal{L}_{AR}(x) = -\sum_{m=1}^{M} \sum_{j=1}^{|\mathcal{T}_m|} \log p_{\theta} (t_{m,j} \mid \mathcal{T}_{< m}, \mathcal{A}_{< m}, t_{m, < j})$$
 (2)

3.2 Absorbing discrete diffusion for audio spans

Building on the theoretical foundation established in Section 2, we apply absorbing discrete diffusion to audio spans $\mathcal{A}_{\leq M} = \cup_{m=1}^M \mathcal{A}_m$. This design choice aligns with the fundamental difference in dependency structures: audio tokens exhibit strong **source**—**target** dependencies (conditioning on source text), making them well-suited for the any-order AR nature of diffusion, while text tokens follow **target**—**target** causal dependencies, better handled by standard AR modeling.

Audio-specific Corruption and Denoising For each training sample, we sample a masking level $\lambda \sim U([0,1])$ and independently mask each audio token with probability λ , while preserving all text tokens. This creates corrupted sequences where audio spans contain a mixture of original tokens and mask tokens $[\mathbf{M}]$, but text spans remain intact. To enable efficient parallel training across all audio spans simultaneously, we apply masking operations to every audio span \mathcal{A}_m in the sequence, rather than processing them sequentially. This parallel masking strategy significantly improves training efficiency while leveraging the time-independent nature of the denoising objective (Eq. 1).

Training Objective for Audio Generation The model learns to predict the original audio tokens for masked positions by minimizing the λ -denoising cross-entropy loss over all audio spans. As discussed in (Ou et al., 2024), this objective is mathematically equivalent to the any-order AR objective, and can be equivalently expressed in the AO-ARM form:

$$\mathcal{L}_{AO}(x) = \sum_{m=1}^{M} \mathbb{E}_{\pi_m \sim U_{\pi_m}} \sum_{j=1}^{|\mathcal{A}_m|} -\log q_{\theta}(a_{m,\pi_m(j)}|\mathcal{T}_{\leq m}, \mathcal{A}_{< m}, a_{m,\pi_m(< j)})$$
(3)

where π_m is a random permutation over the positions within audio span \mathcal{A}_m , and $a_{m,\pi_m(< j)}$ denotes the audio tokens that appear before position j in the permuted order. This formulation makes explicit that the audio generation objective is learning to predict each audio token conditioned on an arbitrary subset of other tokens within the same span, plus the full cross-modal context from text. This any-order AR nature is what enables parallel generation during inference.

3.3 MULTIMODAL FACTORIZATION AND UNIFIED OBJECTIVE

Having established AR modeling for text in Section 3.1 and discrete diffusion for audio in Section 3.2, we now formalize how these two paradigms can be unified within a single probabilistic framework. The key insight is to leverage the distinct dependency structures of each modality through a partial-order factorization that respects the causal nature of text while allowing flexible ordering within audio spans. Recall that text tokens exhibit strong target-target dependencies requiring causal ordering, while audio tokens primarily depend on source-target relationships with their corresponding text. This suggests that within each audio span \mathcal{A}_m , the tokens can be generated in any order as long as they condition on the appropriate cross-modal context $\mathcal{T}_{\leq m} \cup \mathcal{A}_{< m}$. We formalize this intuition using partial orders over token positions.

A partial order on a set V is a binary relation \leq that is reflexive, antisymmetric, and transitive. A set equipped with such a relation is called a partially ordered set (poset). Two elements $a, b \in V$ are comparable if $a \leq b$ or $b \leq a$; otherwise, they are incomparable. An antichain is a subset of V in which every pair of distinct elements is incomparable — that is, no internal ordering constraints exist among them (Davey & Priestley, 2002).

Partial-order Formulation Let (V, \preceq) be a poset over all token indices in the sequence, where V represents all token positions and \preceq encodes precedence relationships. For our interleaved text-audio setting, we define: (1) Each text token $t_{m,j}$ precedes $t_{m,j+1}$ (maintaining left-to-right causality within text spans). (2) All tokens in span m precede all tokens in span m+1 (maintaining cross-span dependencies). (3) Tokens within each audio span \mathcal{A}_m form an antichain under \preceq (no mandatory

internal ordering), but the model is permitted to condition on previously generated tokens within the same span during training and inference under any linear extension.

For any token i, let $\operatorname{Pa}(i)$ denote its set of predecessors under this partial order. By construction, each audio token $a_{m,j}$ has predecessors $\operatorname{Pa}(a_{m,j}) = \mathcal{T}_{\leq m} \cup \mathcal{A}_{< m}$, while for text tokens $\operatorname{Pa}(t_{m,j}) = \mathcal{T}_{< m} \cup \mathcal{A}_{< m} \cup t_{m,< i}$.

Any linear extension ℓ of the partial order (V, \preceq) induces a valid chain-rule factorization: $p(x) = \prod_{j=1}^{|V|} p(x_{\ell(j)} \mid x_{\operatorname{Pa}(\ell(j))})$. Since audio tokens within each span form an antichain, there are multiple valid linear extensions differing only in the within-span ordering of audio tokens. Rather than committing to a single extension, we can *marginalize* over all possible orderings within audio spans.

Order-marginalized Factorization for Audio Spans For an antichain $S \subseteq V$ (such as tokens within an audio span), we define the *order-marginalized conditional* by averaging over all permutations of S: $\tilde{p}_{\theta}(x_S \mid x_{V \setminus S}) = \mathbb{E}_{\pi \in \mathrm{Perm}(S)} \prod_{j \in S} q_{\theta}(x_{\pi(j)} \mid x_{V \setminus S}, x_{\pi(\langle j)})$, where $q_{\theta}(\cdot \mid \cdot)$ represents the any-order AR learned through discrete diffusion. When applied to our audio spans, this gives:

$$\tilde{p}_{\theta}(\mathcal{A}_m \mid \mathcal{T}_{\leq m}, \mathcal{A}_{\leq m}) = \mathbb{E}_{\pi_m \sim U_{\pi_m}} \prod_{j=1}^{|\mathcal{A}_m|} q_{\theta} \left(a_{m, \pi_m(j)} \mid \mathcal{T}_{\leq m}, \mathcal{A}_{\leq m}, a_{m, \pi_m(\leq j)} \right) \tag{4}$$

Intuitively, this averages the likelihood over all possible within-span orderings, reflecting the fact that audio tokens can be generated in any order given the appropriate cross-modal context. Note that while tokens within \mathcal{A}_m form an antichain under the partial order (i.e., no mandatory sequential constraints), the order-marginalized conditional in Eq. 4 allows the model to leverage local target-target dependencies that may arise under specific generation orders. This flexibility enables the model to capture useful intra-span structures when beneficial.

Hybrid AR-NAR Joint Distribution Combining fixed-order AR for text with order-marginalized factorization for audio, our model induces the joint scoring function:

$$\tilde{p}_{\theta}(x) = \prod_{m=1}^{M} \left[\prod_{j=1}^{|\mathcal{T}_{m}|} p_{\theta}(t_{m,j} \mid \mathcal{T}_{\leq m}, \mathcal{A}_{\leq m}, t_{m, \leq j}) \cdot \underbrace{\tilde{p}_{\theta}(\mathcal{A}_{m} \mid \mathcal{T}_{\leq m}, \mathcal{A}_{\leq m})}_{\text{order-marginalized any-order AR for audio}} \right]$$
(5)

This formulation reveals that both modalities are fundamentally AR: text uses a single linear extension (left-to-right), while audio integrates over all linear extensions consistent with the partial order.

Training Objective and Upper Bound Analysis In practice, we cannot directly optimize $\tilde{p}_{\theta}(x)$ because the order-marginalized conditional in Eq. 4 requires computing expectations over all permutations. Instead, we use the training objectives $\mathcal{L}_{AR}(x)$ and $\mathcal{L}_{AO}(x)$ derived in Section 3.2. The key theoretical insight is that our combined training objective provides a tight upper bound on the negative log-likelihood of the desired joint distribution. To see this, consider the audio term:

$$\mathbb{E}_{\pi_{m} \sim U_{\pi_{m}}} \sum_{j=1}^{|\mathcal{A}_{m}|} \left[-\log q_{\theta} \left(a_{m,\pi_{m}(j)} \mid \mathcal{T}_{\leq m}, \mathcal{A}_{< m}, a_{m,\pi_{m}(< j)} \right) \right] \\
\geq -\log \mathbb{E}_{\pi_{m} \sim U_{\pi_{m}}} \prod_{j=1}^{|\mathcal{A}_{m}|} q_{\theta} \left(a_{m,\pi_{m}(j)} \mid \mathcal{T}_{\leq m}, \mathcal{A}_{< m}, a_{m,\pi_{m}(< j)} \right)$$
(6)

The right-hand side is precisely $-\log \tilde{p}_{\theta}(\mathcal{A}_m \mid \mathcal{T}_{\leq m}, \mathcal{A}_{< m})$ from Eq. 4. The left-hand side is exactly the audio loss term for span m in our practical training objective $\mathcal{L}_{AO}(x)$.

To establish the unified upper bound, we now sum the inequality in Eq. 6 over all audio spans $m = 1, \dots, M$:

$$\sum_{m=1}^{M} \mathbb{E}_{\pi_m} \sum_{j=1}^{|\mathcal{A}_m|} \left[-\log q_{\theta} \left(a_{m,\pi_m(j)} \mid \mathcal{T}_{\leq m}, \mathcal{A}_{< m}, a_{m,\pi_m(< j)} \right) \right] \geq \sum_{m=1}^{M} \left(-\log \tilde{p}_{\theta} (\mathcal{A}_m \mid \mathcal{T}_{\leq m}, \mathcal{A}_{< m}) \right)$$

$$(7)$$

The left-hand side is exactly \mathcal{L}_{AO} . For the text terms, the \mathcal{L}_{AR} loss is defined in Eq. 2. Combining the text and audio terms according to the joint factorization in Eq. 5 yields:

$$\mathcal{L}_{\text{Unified}}(x) \triangleq \mathcal{L}_{AR}(x) + \mathcal{L}_{AO}(x) \ge -\log \tilde{p}_{\theta}(x) \tag{8}$$

a detailed derivation of this inequality is provided in Appendix A.1.1, where this final inequality follows from combining the text equality with the audio inequality derived above. Thus, minimizing our practical training objective $\mathcal{L}_{\text{Unified}}(x)$ corresponds to minimizing an upper bound on the negative log-likelihood of the theoretically motivated joint distribution $\tilde{p}_{\theta}(x)$. This result is significant because: (1) It provides theoretical justification for our hybrid AR-NAR training approach. (2) It guarantees that optimizing the computationally tractable objective $\mathcal{L}_{\text{Unified}}(x)$ will not deviate arbitrarily from the theoretically optimal objective $-\log \tilde{p}_{\theta}(x)$.

Training Pipeline and Loss Computation Our training pipeline starts from a pretrained text LLM and expands its vocabulary with discrete audio codebook tokens and control symbols ($\langle SOA \rangle$, $\langle EOA \rangle$). Each training sequence is organized as interleaved text spans and audio spans. We provide an illustration of loss computation in Appendix A.5. Despite its theoretical and practical advantages, the hybrid AR-NAR paradigm introduces a significant train-test discrepancies that can degrade generation quality. During training, audio spans are partially masked according to the diffusion process, while during inference, the model must generate audio and text tokens conditioned on complete text context and previously generated clean audio tokens. To bridge this gap, we propose three principled training strategies:

- Batchwise AR & NAR Objective Mixing (BANOM): With probability $p_{\rm mix}$, we skip the diffusion noise addition process for certain samples and compute gradients only on text tokens using AR loss. This ensures that during training, text tokens occasionally observe clean, unmasked audio spans—matching the inference scenario where text generation conditions on previously generated complete audio content rather than partially masked spans.
- Prefix Preservation Masking (PPM): For a fraction p_{prefix} of training samples, we randomly select a cutoff index m and ensure that all preceding audio spans $\mathcal{A}_{< m} = \{\mathcal{A}_1, \dots, \mathcal{A}_{m-1}\}$ remain unmasked, while applying NAR diffusion loss only to spans $\mathcal{A}_{\geq m} = \{\mathcal{A}_m, \mathcal{A}_{m+1}, \dots, \mathcal{A}_M\}$. This strategy ensures that during training, when generating span \mathcal{A}_m , the model observes clean representations of all previous spans $\mathcal{A}_{< m}$, matching the inference scenario where audio spans are generated sequentially and each span \mathcal{A}_m conditions on fully generated, clean preceding spans $\mathcal{A}_{< m}$ rather than their corrupted, partially masked versions.
- Stochastic Span Truncation (SST): We address the positional bias in $\langle EOA \rangle$ prediction by randomly truncating audio span \mathcal{A}_M during training. Due to disparate tokenization rates between text and audio, audio tokens significantly outnumber text tokens, resulting in fixed-size spans $\mathcal{A}_1, \ldots, \mathcal{A}_{M-1}$ and a variable-length final span \mathcal{A}_M . Since all audio spans undergo simultaneous diffusion training, the model learns to predict $\langle EOA \rangle$ at fixed positions for early spans, creating a strong positional bias that hinders content-aware termination learning for the final span. To mitigate this, we implement stochastic truncation: with probability p_{trunc} , we randomly select a truncation length $k < |\mathcal{A}_M|$ and create a truncated span $\mathcal{A}_M^{\text{trunc}} = (a_{M,1}, \ldots, a_{M,k})$ by removing the original $\langle EOA \rangle$ token and suffix tokens $(a_{M,k+1}, \ldots, a_{M,|\mathcal{A}_M|})$. This creates training samples where span termination occurs at arbitrary positions rather than fixed boundaries, forcing the model to predict $\langle EOA \rangle$ based on semantic content and contextual text rather than positional cues.

3.4 MODALITY-AWARE ATTENTION MECHANISM

Our attention design enforces a step-wise pattern across three content types: (1) the input prompt uses standard causal attention; (2) text tokens \mathcal{T}_m apply strict causal attention to the prompt, all

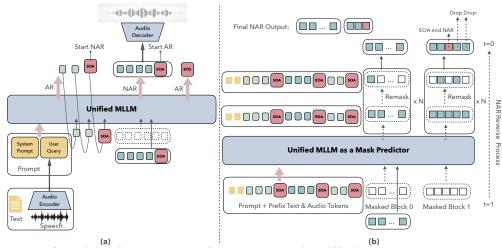


Figure 2: Overview of the proposed framework and its diffusion reverse process. (a) TtT framework. A unified MLLM that interleaves AR text and NAR audio generation. The model alternates between AR text decoding and NAR audio synthesis based on control tokens. (b) Diffusion reverse process. NAR audio generation through iterative denoising

prior spans, and preceding tokens within their span; and (3) audio tokens A_m use hybrid attention—bidirectional within their own span and causal attention to the prompt and all earlier spans. This pattern enables parallel audio span training in a single forward pass while preventing cross-span interference. See Appendix A.5 for an illustration.

3.5 Inference Process

Figure 2 shows the overview of TtT and its inference process. During inference, TtT alternates between AR text decoding and NAR audio synthesis within a unified framework. Given input audio, the encoder produces discrete tokens, which are processed by AR generation until $\langle SOA \rangle$ is reached. The model then switches to NAR mode, where block-wise diffusion (see Algorithm 1 in Appendix A.4) generates audio spans in parallel. Upon predicting $\langle EOA \rangle$, remaining tokens in the block are dropped and decoding returns to AR mode, repeating the cycle until $\langle EOS \rangle$ is generated. Each completed audio span is immediately sent to the audio decoder, enabling parallel synthesis with low first-token latency and continuous streaming generation.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUPS

Datasets To effectively train and evaluate our proposed TtT framework, we follow prior works (Zeng et al., 2025; Ding et al., 2025; Long et al., 2025) and adopt a diverse collection of multi-task datasets, including ASR, TTS, audio chat, text chat, Automated Audio Captioning (AAC), Speech Emotion Classification (SEC), Acoustic Scene Classification (ASC), and interleaved text–audio data. In total, the combined training corpus contains approximately 6.3 million samples across these tasks. For evaluation, we focus on two representative tasks: Audio-QA, which evaluates both the model's reasoning skills and its text-to-speech capabilities in an end-to-end audio generation framework, and ASR, which measures the quality of cross-modal alignment between text and audio. We provide details in Appendix A.6

Evaluation To effectively evaluate our model on two representative tasks, we carefully design the evaluation protocol and metrics. For Audio-QA, we introduce an ASR-LLM pipeline that transcribes the model's spoken responses using language-specific ASR systems (Paraformer-zh for Chinese and Whisper-Large-v3 for English) and leverages a powerful LLM-as-a-Judge (Qwen3-235B-A30B) to assess semantic correctness against ground-truth answers. This approach captures the model's ability to generate natural, accurate spoken responses in realistic conversational settings. For the

Table 1: Comprehensive evaluation of TtT framework: performance comparison, ablation study, and training strategy analysis. Higher (\uparrow) is better for Audio-QA, lower (\downarrow) is better for ASR. Datasets abbreviations are available in Table 4

Models	Audio-QA (↑)			ASR (↓)						
1,104015	AE.	LQ.	TQA.	WQ.	Fzh.	A2.	A1.	WS_met.	WS_net.	Fen.
	Main Results									
Qwen2.5-1.5B (AR)	17.99	16.78	1.61	2.32	99.08	59.73	80.27	85.55	81.76	96.16
Qwen2.5-1.5B (NAR)	10.70	0.00	0.40	0.20	86.97	224.37	191.11	123.96	143.76	108.25
TtT-1.5B (AR-NAR)	15.68	23.75	3.47	7.70	44.36	14.89	16.72	52.23	41.52	49.00
Qwen2.5-3B (AR)	14.42	10.00	0.60	0.70	90.32	54.94	72.01	80.01	73.64	74.47
Qwen2.5-3B (NAR)	11.31	0.67	1.21	0.70	68.94	212.27	160.58	89.22	111.29	83.51
TtT-3B (AR-NAR)	17.46	34.68	6.53	11.61	55.67	12.53	13.65	53.83	44.29	64.31
Ablation Study										
TtT-3B w/o BANOM	13.87	19.87	2.81	5.12	58.25	18.58	21.35	58.48	49.52	68.90
TtT-3B w/o PPM	14.27	22.79	2.71	5.54	58.86	15.63	18.83	57.76	47.92	67.37
TtT-3B w/o SST	14.12	10.20	1.30	3.72	56.39	25.43	31.03	64.41	56.70	62.60
TtT-3B (AR-NAR)	17.46	34.68	6.53	11.61	55.67	12.53	13.65	53.83	44.29	64.31
Training Strategy Comparison										
TtT-3B (AR-NAR)	17.46	34.68	6.53	11.61	55.67	12.53	13.65	53.83	44.29	64.31
Pretrain+AR	29.45	15.93	3.61	11.45	23.37	9.79	12.67	26.75	20.91	19.49
Pretrain+TtT	26.73	40.07	11.07	21.43	18.99	6.80	5.78	27.59	19.85	19.10

ASR task, we directly measure transcription accuracy using Word Error Rate (WER), which reflects both speech recognition fidelity and cross-modal alignment within our hybrid AR-NAR framework. More details are in Appendix A.7.1.

Model Configuration We adopt the Qwen2.5-Base model as the backbone, experimenting with parameter scales of 1.5B and 3B, and fine-tune all parameters during training. For the audio components, we directly follow the audio tokenizer and decoder design introduced in GLM-4-Voice (Zeng et al., 2024). These modules have been shown to provide efficient and high-quality speech tokenization and synthesis, and they allow our framework to leverage strong audio modeling without requiring additional architectural modifications. The training details are provide in Appendix A.8.

4.2 MAIN RESULTS

To comprehensively evaluate the effectiveness of our proposed TtT framework, we compare it with two representative baselines, one employing a purely AR backbone and the other adopting a purely diffusion based NAR backbone. For fairness and scalability, all three frameworks are instantiated with backbones of 1.5B and 3B parameters.

Conversational Understanding Ability Table 1 main results part reports the comparison of conversational understanding ability across different models on Audio-QA task. From these results we make the following observations: (1) Our proposed TtT consistently achieves the best performance on all four Audio-QA datasets. For example, at the 3B scale, TtT-3B attains scores of 17.46, 34.68, 6.53, and 11.61 on AlpacaEval, LLaMAQuestions, TriviaQA, and WebQuestions respectively, surpassing the Qwen2.5-3B (AR) model by +3.04, +24.68, +5.93, and +10.91. This improvement stems from the hybrid AR–NAR design, which better captures the inherent asymmetry between text and audio dependencies. (2) Our results demonstrate a clear scaling trend shared with the baselines, where larger backbones consistently yield stronger performance. For instance, TtT-3B achieves 34.68 on LLaMAQuestions, outperforming its 1.5B counterpart which scores 23.75, reflecting an absolute gain of +10.93. (3) Qwen2.5-1.5B (NAR) and Qwen2.5-3B (NAR) perform noticeably worse than both our hybrid framework and the AR-only baseline. Since discrete diffusion models employ an order-agnostic AO-ARM objective that ignore the inherent sequential structure of interleaved text-audio sequences. This mismatch leads to order confusion, compromising sequence coherence and weakening audio-text alignment.

Cross-Modal Alignment Ability — As shown in Table 1 main results part, our proposed TtT significantly outperforms both AR and NAR on ASR tasks, demonstrating superior cross-modal alignment. For example, at the 3B scale, TtT-3B achieves 12.53 and 13.65 WER on AISHELL-2 and AISHELL-1 respectively, yielding improvements of 42.41 and 58.36 absolute WER points over Qwen2.5-3B (AR). This advantage comes from the hybrid AR–NAR design, where NAR diffusion enables efficient parallel denoising for tighter audio—text alignment, and AR components maintain coherent cross-modal conditioning. Similar to the conversational understanding task, we observe a clear scaling trend in cross-modal alignment. Larger backbones consistently deliver better alignment, with TtT-3B achieving substantially lower WER than TtT-1.5B (e.g., 12.53 vs. 14.89 on AISHELL-2, and 13.65 vs. 16.72 on AISHELL-1).

4.3 ABLATION STUDY

To better understand the contribution of each training strategy in our hybrid AR-NAR framework, we perform an ablation study based on the full model TtT-3B (AR-NAR). The variant w/o BANOM corresponds to removing batchwise AR & NAR objective mixing from the full model, w/o PPM removes prefix preservation masking, and w/o SST removes stochastic span truncation. Table 1 ablation study part presents the detailed results of our ablation experiments. From these results we draw the following conclusions: (1) All three training strategies have a positive impact on model performance, and removing any one of them leads to clear degradation. For instance, on the LLa-MAQuestions dataset, removing SST reduces the score from 34.68 to 10.20. This drop occurs because stochastic truncation mitigates positional bias in $\langle EOA \rangle$ prediction, forcing span termination by semantic content rather than position. Removing it weakens variable-length audio generation and reduces flexibility in conversational outputs. (2) Removing BANOM yields the largest performance degradation. For example, on the AISHELL-2 dataset, the performance decreases from 12.53 to 18.58 when the strategy is removed. This mechanism is essential for exposing text tokens to clean audio prefixes during training, better matching inference. Without it, the model faces sharper train—test discrepancy, weakening cross-modal consistency and alignment.

4.4 EFFECT OF MULTIMODAL ALIGNMENT PRETRAINING.

To further investigate the effectiveness of our method on top of a multimodally aligned pretrained model, we perform large-scale multimodal pretraining based on the Qwen2.5-3B-Base model. Specifically, we construct a corpus of approximately 200B tokens covering ASR, TTS, text-only data, and interleaved text-audio data. The model is trained with a standard AR objective using a global batch size of 256 for 140k steps. This pretraining stage equips the backbone model (Qwen2.5-3B-Base) with strong cross-modal alignment ability before applying our hybrid AR-NAR learning framework. Table 1 compares the AR-only and AR-NAR frameworks under two different training strategies, specifically training directly from Qwen2.5-3B-Base without multimodal pretraining (TtT-3B) and initialization from the multimodally aligned pretrained model (Pretrain+AR and Pretrain+TtT). From the table, we observe that: (1) When trained directly from Qwen2.5-3B-Base, our TtT framework achieves comparable or even superior performance to the AR-only baseline, indicating that the hybrid AR-NAR design is already competitive without pretraining; (2) when applied on top of the multimodally aligned pretrained model, Pretrain+TtT consistently matches or surpasses Pretrain+AR across both Audio-QA and ASR tasks. These results demonstrate that TtT not only performs strongly from scratch, but also benefits significantly when built upon large-scale multimodal alignment pretraining.

5 CONCLUSION

In this work, we introduce a unified framework that combines autoregressive text generation with non-autoregressive audio diffusion. By explicitly respecting the asymmetry between text and audio dependencies, our framework bridges the strengths of AR and NAR modeling within a single Transformer. We further propose simple yet effective strategies to mitigate train–test discrepancies, enabling robust and flexible audio generation. Experiments on Audio-QA and ASR benchmarks demonstrate clear improvements over strong AR and NAR baselines. Our results highlight the importance of modality-aware design for building scalable and effective speech-to-speech systems.

REPRODUCIBILITY STATEMENT

The anonymous downloadable source code is available at: https://anonymous.4open.science/r/TtT. For theoretical results, a complete proof of the claim in included in the Section A.1.1 in Appendix.

REFERENCES

- Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne Van Den Berg. Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems*, 34:17981–17993, 2021.
- Zalán Borsos, Raphaël Marinier, Damien Vincent, Eugene Kharitonov, Olivier Pietquin, Matt Sharifi, Dominik Roblek, Olivier Teboul, David Grangier, Marco Tagliasacchi, et al. Audiolm: a language modeling approach to audio generation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31:2523–2533, 2023.
- George EP Box, Gwilym M Jenkins, Gregory C Reinsel, and Greta M Ljung. *Time series analysis: forecasting and control.* John Wiley & Sons, 2015.
- Junyi Chen, Shuming Shen, Andi Chen, Wen Wu, Jiantao Kang, Haohe Li, Weijiang Zhou, Yi Ren, Yanmin Qian, Xipeng Qiu, et al. Speecht5: Unified-modal encoder-decoder pre-training for spoken language processing. *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, 2022.
- Yunfei Chu, Jin Xu, Qian Yang, Haojie Wei, Xipin Wei, Zhifang Guo, Yichong Leng, Yuanjun Lv, Jinzheng He, Junyang Lin, et al. Qwen2-audio technical report. *arXiv preprint arXiv:2407.10759*, 2024.
- Trung Dang, David Aponte, Dung Tran, and Kazuhito Koishida. Livespeech: Low-latency zero-shot text-to-speech via autoregressive modeling of audio discrete codes. *arXiv preprint arXiv:2406.02897*, 2024.
- Brian A Davey and Hilary A Priestley. *Introduction to lattices and order*. Cambridge university press, 2002.
- Alexandre Défossez, Laurent Mazaré, Manu Orsini, Amélie Royer, Patrick Pérez, Hervé Jégou, Edouard Grave, and Neil Zeghidour. Moshi: a speech-text foundation model for real-time dialogue. arXiv preprint arXiv:2410.00037, 2024.
- Ding Ding, Zeqian Ju, Yichong Leng, Songxiang Liu, Tong Liu, Zeyu Shang, Kai Shen, Wei Song, Xu Tan, Heyi Tang, et al. Kimi-audio technical report. *arXiv preprint arXiv:2504.18425*, 2025.
- Guhao Feng, Yihan Geng, Jian Guan, Wei Wu, Liwei Wang, and Di He. Theoretical benefit and limitation of diffusion language model. *arXiv preprint arXiv:2502.09622*, 2025.
- Shansan Gong, Shivam Agarwal, Yizhe Zhang, Jiacheng Ye, Lin Zheng, Mukai Li, Chenxin An, Peilin Zhao, Wei Bi, Jiawei Han, et al. Scaling diffusion language models via adaptation from autoregressive models. *arXiv preprint arXiv:2410.17891*, 2024.
- Ailin Huang, Boyong Wu, Bruce Wang, Chao Yan, Chen Hu, Chengli Feng, Fei Tian, Feiyu Shen, Jingbei Li, Mingrui Chen, et al. Step-audio: Unified understanding and generation in intelligent speech interaction. *arXiv preprint arXiv:2502.11946*, 2025.
- Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae. Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis. *Advances in Neural Information Processing Systems*, 33:17022–17033, 2020.
 - Gen Li and Changxiao Cai. A convergence theory for diffusion language models: An information-theoretic perspective. *arXiv preprint arXiv:2505.21400*, 2025.

Tianpeng Li, Jun Liu, Tao Zhang, Yuanbo Fang, Da Pan, Mingrui Wang, Zheng Liang, Zehuan Li, Mingan Lin, Guosheng Dong, et al. Baichuan-audio: A unified framework for end-to-end speech interaction. *arXiv* preprint arXiv:2502.17239, 2025.

- Alexander H Liu, Sang-gil Lee, Chao-Han Huck Yang, Yuan Gong, Yu-Chiang Frank Wang, James R Glass, Rafael Valle, and Bryan Catanzaro. Uniwav: Towards unified pre-training for speech representation learning and generation. *arXiv preprint arXiv:2503.00733*, 2025.
- Zuwei Long, Yunhang Shen, Chaoyou Fu, Heting Gao, Lijiang Li, Peixian Chen, Mengdan Zhang, Hang Shao, Jian Li, Jinlong Peng, et al. Vita-audio: Fast interleaved cross-modal token generation for efficient large speech-language model. *arXiv preprint arXiv:2505.03739*, 2025.
- Justin Lovelace, Varsha Kishore, Yiwei Chen, and Kilian Weinberger. Diffusion guided language modeling. *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 14936–14952, 2024.
- Shivam Mehta, Ruibo Tu, Jonas Beskow, Éva Székely, and Gustav Eje Henter. Matcha-tts: A fast tts architecture with conditional flow matching. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 11341–11345. IEEE, 2024.
- Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin, Ji-Rong Wen, and Chongxuan Li. Large language diffusion models. *arXiv preprint arXiv:2502.09992*, 2025.
- Jingyang Ou, Shen Nie, Kaiwen Xue, Fengqi Zhu, Jiacheng Sun, Zhenguo Li, and Chongxuan Li. Your absorbing discrete diffusion secretly models the conditional distributions of clean data. arXiv preprint arXiv:2406.03736, 2024.
- Marc' Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. Sequence level training with recurrent neural networks. *arXiv preprint arXiv:1511.06732*, 2015.
- Yi Ren, Jinglin Liu, Xu Tan, Zhou Zhao, Sheng Zhao, and Tie-Yan Liu. A study of non-autoregressive model for sequence generation. *arXiv preprint arXiv:2004.10454*, 2020.
- Paul K Rubenstein, Chulayuth Asawaroengchai, Duc Dung Nguyen, Ankur Bapna, Zalán Borsos, Félix de Chaumont Quitry, Peter Chen, Dalia El Badawy, Wei Han, Eugene Kharitonov, et al. Audiopalm: A large language model that can speak and listen. *arXiv preprint arXiv:2306.12925*, 2023.
- Subham Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin Chiu, Alexander Rush, and Volodymyr Kuleshov. Simple and effective masked diffusion language models. *Advances in Neural Information Processing Systems*, 37:130136–130184, 2024.
- Jiaxin Shi, Kehang Han, Zhe Wang, Arnaud Doucet, and Michalis Titsias. Simplified and generalized masked diffusion for discrete data. *Advances in Neural Information Processing Systems*, 37: 103131–103167, 2024.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- Chengyi Wang, Sanyuan Chen, Yu Wu, Ziyang Zhang, Long Zhou, Shujie Liu, Zhuo Chen, Yanqing Liu, Huaming Wang, Jinyu Li, et al. Neural codec language models are zero-shot text to speech synthesizers. *arXiv preprint arXiv:2301.02111*, 2023.
- Zhifei Xie and Changqiao Wu. Mini-omni: Language models can hear, talk while thinking in streaming. *arXiv preprint arXiv:2408.16725*, 2024.
 - Jin Xu, Zhifang Guo, Jinzheng He, Hangrui Hu, Ting He, Shuai Bai, Keqin Chen, Jialin Wang, Yang Fan, Kai Dang, et al. Qwen2. 5-omni technical report. *arXiv preprint arXiv:2503.20215*, 2025.
- Ling Yang, Ye Tian, Bowen Li, Xinchen Zhang, Ke Shen, Yunhai Tong, and Mengdi Wang. Mmada: Multimodal large diffusion language models. *arXiv preprint arXiv:2505.15809*, 2025.

Jiacheng Ye, Jiahui Gao, Shansan Gong, Lin Zheng, Xin Jiang, Zhenguo Li, and Lingpeng Kong. Beyond autoregression: Discrete diffusion for complex reasoning and planning. *arXiv preprint arXiv:2410.14157*, 2024.

- Runpeng Yu, Xinyin Ma, and Xinchao Wang. Dimple: Discrete diffusion multimodal large language model with parallel decoding. *arXiv preprint arXiv:2505.16990*, 2025.
- Aohan Zeng, Zhengxiao Du, Mingdao Liu, Kedong Wang, Shengmin Jiang, Lei Zhao, Yuxiao Dong, and Jie Tang. Glm-4-voice: Towards intelligent and human-like end-to-end spoken chatbot. *arXiv* preprint arXiv:2412.02612, 2024.
- Aohan Zeng, Zhengxiao Du, Mingdao Liu, Lei Zhang, Shengmin Jiang, Yuxiao Dong, and Jie Tang. Scaling speech-text pre-training with synthetic interleaved data. In *Proceedings of the 13th International Conference on Learning Representations*, Singapore, April 2025.

A APPENDIX

A.1 MATHEMATICAL DERIVATION

A.1.1 DERIVATION OF THE TRAINING OBJECTIVE UPPER BOUND

Recall from Eq. 5 that the joint distribution factors as:

$$\tilde{p}_{\theta}(x) = \prod_{m=1}^{M} \left[\prod_{j=1}^{|\mathcal{T}_m|} p_{\theta}(t_{m,j} \mid \mathcal{T}_{\leq m}, \mathcal{A}_{\leq m}, t_{m,\leq j}) \cdot \tilde{p}_{\theta}(\mathcal{A}_m \mid \mathcal{T}_{\leq m}, \mathcal{A}_{\leq m}) \right]. \tag{9}$$

Taking the negative logarithm of both sides gives:

$$-\log \tilde{p}_{\theta}(x) = -\sum_{m=1}^{M} \sum_{j=1}^{|\mathcal{T}_{m}|} \log p_{\theta}(t_{m,j} \mid \mathcal{T}_{\leq m}, \mathcal{A}_{\leq m}, t_{m,\leq j}) - \sum_{m=1}^{M} \log \tilde{p}_{\theta}(\mathcal{A}_{m} \mid \mathcal{T}_{\leq m}, \mathcal{A}_{\leq m})$$

$$= \mathcal{L}_{AR}(x) + \sum_{m=1}^{M} \left(-\log \tilde{p}_{\theta}(\mathcal{A}_{m} \mid \mathcal{T}_{\leq m}, \mathcal{A}_{\leq m}) \right). \tag{10}$$

By Eq. 6 and its summation over m, we have:

$$\mathcal{L}_{AO}(x) \ge \sum_{m=1}^{M} \left(-\log \tilde{p}_{\theta}(\mathcal{A}_m \mid \mathcal{T}_{\le m}, \mathcal{A}_{\le m}) \right). \tag{11}$$

Therefore, combining both components:

$$\mathcal{L}_{AR}(x) + \mathcal{L}_{AO}(x) \ge \mathcal{L}_{AR}(x) + \sum_{m=1}^{M} \left(-\log \tilde{p}_{\theta}(\mathcal{A}_m \mid \mathcal{T}_{\le m}, \mathcal{A}_{\le m}) \right) = -\log \tilde{p}_{\theta}(x), \quad (12)$$

which establishes Eq. 8. This confirms that our practical training objective $\mathcal{L}_{\text{Unified}}(x)$ is a valid upper bound on the true negative log-likelihood, enabling tractable optimization while preserving consistency with the target joint distribution $\tilde{p}_{\theta}(x)$.

A.2 RELATED WORK

A.2.1 AUDIO-LANGUAGE MODEL PRETRAINING

Recent advances in end-to-end audio-language models have moved beyond traditional cascaded architectures (Chen et al., 2022; Wang et al., 2023) toward unified multimodal frameworks. Representative works include Moshi (Défossez et al., 2024), which achieves real-time duplex speech conversation through hierarchical Transformer architectures; GLM4-Voice (Zeng et al., 2024), which builds upon GLM-4-9B for robust Chinese and English speech processing; and VITA-Audio (Long et al., 2025), which introduces a lightweight Multiple Cross-modal Token Prediction (MCTP) module for fast audio-text generation with significantly reduced first-token latency. More recent efforts have focused on scaling and production readiness: Step-Audio (Huang et al., 2025) presents a 130B-parameter unified speech-text model with generative speech data engine and instruction-driven fine control across dialects, emotions, singing, and RAP, while Baichuan-Audio (Li et al., 2025) features text-guided aligned speech generation with multi-codebook discretization to preserve both semantic and acoustic information. UniWav (Liu et al., 2025) proposes the first unified encoder-decoder framework that jointly learns representation encoders and generative audio decoders for both discriminative and generative speech tasks.

A key limitation shared by these approaches is their reliance on uniform autoregressive objectives for both text and audio tokens, which overlooks the distinct dependency structures of these modalities. Our work addresses this gap by proposing a hybrid AR-NAR framework that respects the inherent asymmetries between text and audio generation.

A.2.2 DISCRETE DIFFUSION MODELS

Discrete diffusion models have emerged as a compelling alternative to autoregressive generation, offering non-autoregressive approaches that can generate entire sequences in parallel. The foundational work of D3PMs (Austin et al., 2021) generalized diffusion processes to discrete data through flexible transition matrices, with absorbing processes that progressively mask tokens proving particularly effective. This framework has since evolved through both theoretical advances and practical improvements. From a theoretical perspective, recent work has deepened our understanding of discrete diffusion dynamics. Ou et al. (2024) revealed that absorbing diffusion's concrete score can be expressed as time-independent conditional probabilities, leading to RADD—a reparameterized model that removes explicit time conditioning while establishing connections to any-order autoregressive generation. Building on this foundation, Li & Cai (2025) formally characterized convergence rates, proving that KL divergence decays at O(1/T) with bounds scaling linearly with token mutual information. However, Feng et al. (2025) identified a fundamental trade-off: while masked diffusion achieves near-optimal perplexity in constant steps, sequence-level tasks like reasoning may require steps linear in sequence length. Practical advances have focused on training efficiency and application domains. Shi et al. (2024) reformulated the variational objective as a weighted integral of cross-entropy losses, unifying prior approaches while achieving state-of-the-art results that even surpass comparable autoregressive baselines. For complex reasoning tasks where autoregressive models struggle with subgoal imbalance, Ye et al. (2024) demonstrated that Multi-Granularity Diffusion Modeling can achieve near-perfect accuracy by prioritizing harder subgoals during training. The scalability challenge has been addressed through innovative adaptation strategies. Rather than training from scratch, Gong et al. (2024); Nie et al. (2025) showed that pretrained autoregressive models can be efficiently converted to diffusion models via continual pre-training, maintaining competitive performance while enabling parallel generation. Meanwhile, hybrid approaches are gaining traction: Lovelace et al. (2024) combined diffusion-based latent proposals with autoregressive decoding for controllable generation, while Yang et al. (2025) developed MMaDA, a unified multimodal diffusion foundation model that processes text, images, and reasoning within a single architecture.

A.3 AUTOREGRESSIVE MODELING & ABSORBING DISCRETE DIFFUSION

A.3.1 AUTOREGRESSIVE MODELING

Autoregressive (AR) models are a fundamental class of generative models that factorize the joint probability distribution of a sequence $x=(x^1,\ldots,x^L)$ into a product of conditional probabilities, based on the chain rule:

$$p(x) = \prod_{i=1}^{L} p(x^{i}|x^{< i})$$
(13)

where $x^{< i} = (x^1, \dots, x^{i-1})$ represents the tokens preceding the current token x^i . This factorization imposes a sequential, causal structure on the generation process. Such models, typically implemented with Transformer decoders, are trained by minimizing the negative log-likelihood (NLL) of the data, which corresponds to a cross-entropy loss at each position.

A.3.2 ABSORBING DISCRETE DIFFUSION

Discrete diffusion models offer a non-autoregressive alternative for sequence generation. We focus on absorbing discrete diffusion (Austin et al., 2021; Ou et al., 2024), which involves a forward corruption process and a learned reverse denoising process.

Forward Process. The forward process is a continuous-time discrete Markov chain that corrupts a clean sequence x_0 over a time interval $t \in [0,T]$. Its dynamics are governed by a time-dependent transition rate matrix $Q_t = \sigma(t)Q$, where $\sigma(t)$ is a positive noise schedule. For absorbing diffusion, the constant matrix $Q = Q^{\text{abs}}$ is defined as:

$$Q^{\text{abs}}(x \to x') = \begin{cases} 1, & \text{if } x' = [\mathbf{M}] \text{ and } x \neq [\mathbf{M}], \\ -1, & \text{if } x' = x \neq [\mathbf{M}], \\ 0, & \text{otherwise.} \end{cases}$$
(14)

This structure dictates that any token $x \neq [\mathbf{M}]$ transitions to a special mask token $[\mathbf{M}]$ at a rate of $\sigma(t)$. The state $[\mathbf{M}]$ is an **absorbing state** because the transition rate out of it is zero (i.e., $\mathbf{Q}^{\mathrm{abs}}([\mathbf{M}] \to x') = 0$ for all x'), meaning once a token is masked, it remains masked. Over time, the sequence converges to a fully masked state. The probability that a token is masked by time t is given by $\lambda(t) = 1 - e^{-\int_0^t \sigma(s) ds}$.

Reverse Process. The reverse process is also a continuous-time Markov chain that learns to denoise a corrupted sequence x_t back towards the clean data x_0 . Its reverse transition rate matrix \tilde{Q}_t is related to the forward rate matrix by:

$$\tilde{\mathbf{Q}}_t(x_t \to \hat{x}_t) = \begin{cases} \mathbf{Q}_t(\hat{x}_t \to x_t) \frac{p_t(\hat{x}_t)}{p_t(x_t)}, & x_t \neq \hat{x}_t, \\ -\sum_{k \neq x} \tilde{\mathbf{Q}}_t(x_t, k), & \hat{x}_t = x_t. \end{cases}$$
(15)

The term $p_t(\hat{x}_t)/p_t(x_t)$ is known as the *concrete score*. Since the forward process only allows transitions to the $[\mathbf{M}]$ state, the only non-trivial reverse transitions are from $[\mathbf{M}]$ back to a vocabulary token. This simplifies the learning task to modeling the score for these specific denoising transitions.

Time-Independent Score and the Denoising Objective. A key theoretical insight for absorbing diffusion is that the concrete score analytically decomposes into a known, time-dependent scalar and a *time-independent* conditional probability over the clean data (Ou et al., 2024). Specifically, for a transition that unmasks position i from $[\mathbf{M}]$ to a token v, the score is:

$$\underbrace{\frac{p_t(\dots,\hat{x}^i=v,\dots)}{p_t(\dots,x^i=[\mathbf{M}],\dots)}}_{\text{concrete score}} = \underbrace{\frac{e^{-\bar{\sigma}(t)}}{1-e^{-\bar{\sigma}(t)}}}_{\text{time scalar}} \cdot \underbrace{p_0(v \mid UM)}_{\text{clean conditional probability}}.$$
(16)

where UM denotes the set of unmasked (visible) tokens in the corrupted sequence. This decomposition is crucial because it decouples the time-dependent dynamics from the data distribution. It implies that the model q_{θ} does not need to learn a complex function of time t. Instead, its sole objective is to learn to approximate the clean conditional distribution $p_0(v|UM)$, which is a static, time-independent property of the data. The learning task is thus simplified to a denoising objective: given a corrupted sequence with some tokens masked, predict the original tokens for the masked positions based on the visible context.

Equivalence to Any-Order Autoregressive Modeling. This denoising perspective reveals a profound connection to autoregressive modeling. A standard AR model learns to predict a token based on a fixed, causal context. The diffusion model, through its denoising objective, learns to predict a token given an arbitrary context of unmasked tokens. This ability to condition on any subset of the context is the defining feature of an Any-Order Autoregressive Model (AO-ARM).

In fact, the principled training objective for the diffusion model, known as the λ -denoising cross-entropy loss, is mathematically equivalent to the training objective of an AO-ARM (Ou et al., 2024), which averages the prediction loss over all possible permutations (or orderings) of the sequence:

$$\mathcal{L}_{AO}(\boldsymbol{x}_0) = \mathbb{E}_{\pi \sim U_{\pi}} \sum_{l=1}^{d} -\log q_{\theta}(x_0^{\pi(l)} | x_0^{\pi(< l)}). \tag{17}$$

where π is a random permutation of the token indices. Therefore, training an absorbing discrete diffusion model is equivalent to training a powerful ensemble of autoregressive models that can operate in any order. This inherent flexibility is what enables parallel, non-autoregressive generation at inference time and makes it a suitable choice for modeling source-dependent modalities like audio.

A.4 BLOCK-WISE MASKED DIFFUSION GENERATION FOR AUDIO TOKENS

For NAR audio generation, we employ a block-wise denoising approach adapted from Nie et al. (2025). Unlike full-sequence diffusion, it processes audio in fixed-length blocks, balancing parallelism and controllability.

As detailed in Algorithm 1, the model generates audio in fixed-size blocks of length B, where each block is progressively denoised over T steps using an absorbing discrete diffusion process.

At each denoising step t, the model predicts tokens for all currently masked positions in parallel. The algorithm then selectively commits the most confident predictions (determined by predicted probability or random sampling) while remasking the remaining positions for further refinement. This progressive denoising continues until all positions in the current block are decoded. Crucially, if an $\langle \text{EOA} \rangle$ token is generated within a block, decoding terminates immediately at that position, truncating the remainder and seamlessly returning control to the AR text generation mode.

810

811

812

813

814

815 816

817

```
Algorithm 1 Block-wise Masked Diffusion for Autoregressive Audio Generation
```

```
818
               Require: Context tokens \mathbf{c} \in \mathbb{N}^{1 \times L_c}, max generation length L_{\max} \in \mathbb{N},
819
                 1: Sampling steps T \in \mathbb{N}, block length B \in \mathbb{N}, temperature \tau \geq 0,
820
                 2: CFG scale \gamma \geq 0, remasking strategy \mathcal{R} \in \{\text{low\_confidence}, \text{random}\},
821
                 3: Special token IDs: mask m_{\text{mask}}, end-of-audio \mathcal{E}.
822
               Ensure: Generated token sequence \mathbf{s} \in \mathbb{N}^{1 \times L} with L \leq L_c + L_{\text{max}}.
823
                 4: Initialize \mathbf{s} \leftarrow \mathbf{c}

    Start from context

                 5: while |\mathbf{s}| < |\mathbf{c}| + L_{\max} do
824
                            \mathbf{x} \leftarrow \text{pad}(\mathbf{s}, B, \text{value} = m_{\text{mask}})
                                                                                                                                             \triangleright Append B mask tokens
                 6:
825
                            \mathcal{M}_{\text{block}} \leftarrow \{i \mid \mathbf{x}_i = m_{\text{mask}} \land i \geq |\mathbf{s}|\} \\ \{n_t\}_{t=1}^T \leftarrow \text{schedule}(|\mathcal{M}_{\text{block}}|, T)
                 7:
                                                                                                                                                826
                 8:
                                                                                                                                         827
                 9:
                            for t = 1 to T do
828
                                   \mathcal{M}_t \leftarrow \{i \mid \mathbf{x}_i = m_{\text{mask}}\}
               10:
                                                                                                                                              829
                                   if \gamma > 0 then
               11:
830
                                         \mathbf{x}_{\text{uncond}} \leftarrow \mathbf{x}; \, \mathbf{x}_{\text{uncond}}[\neg \mathcal{M}_t] \leftarrow m_{\text{mask}}
               12:
                                                                                                                                                   ▶ Unconditional input
831
                                                                                                                                                         ▶ Batched forward
               13.
                                         \ell_{\text{cond}}, \ell_{\text{uncond}} \leftarrow \text{model}([\mathbf{x}; \mathbf{x}_{\text{uncond}}])
832
                                         \ell \leftarrow \ell_{\text{uncond}} + (\gamma + 1) \cdot (\ell_{\text{cond}} - \ell_{\text{uncond}})
               14:
833
               15:
                                   else
834
               16:
                                         \ell \leftarrow model(\mathbf{x})
               17:
                                   end if
835
                                   \hat{\mathbf{x}} \leftarrow \arg\max(\text{Gumbel}(\boldsymbol{\ell}, \tau))
                                                                                                                                                       18:
836
               19:
                                  if \mathcal{R} = \text{low\_confidence} then
837
                                         \mathbf{p} \leftarrow \text{softmax}(\boldsymbol{\ell}); \mathbf{c}_i \leftarrow \mathbf{p}_i[\hat{\mathbf{x}}_i]
                                                                                                                                    Confidence = predicted prob
               20:
838
                                   else if \mathcal{R} = \text{random then}
               21:
839
               22:
                                         \mathbf{c}_i \leftarrow \mathtt{Uniform}(0,1) \text{ for } i \in \mathcal{M}_t
840
                                   end if
               23:
841
                                   \mathbf{c}_i \leftarrow -\infty \text{ for } i < |\mathbf{s}|
                                                                                                                                                ▶ Protect context tokens
               24:
842
               25:
                                   \hat{\mathbf{x}}_i \leftarrow \mathbf{x}_i \text{ for } i \notin \mathcal{M}_t
                                                                                                                                  ▷ Only update masked positions
843
                                   \mathcal{K}_t \leftarrow \text{TopK}(\{\mathbf{c}_i \mid i \in \mathcal{M}_t\}, k = n_t)
               26:
                                                                                                                \triangleright Select n_t most confident/random tokens
844
               27:
                                   \mathbf{x}_i \leftarrow \hat{\mathbf{x}}_i \text{ for all } i \in \mathcal{K}_t
                                                                                                                                      end for
845
               28:
                            \mathbf{b} \leftarrow \mathbf{x}[|\mathbf{s}| : |\mathbf{s}| + B]
               29:
                                                                                                                                             846
               30:
                            if \mathcal{E} \cap \mathbf{b} \neq \emptyset then
847
                                   p \leftarrow \min\{i \mid \mathbf{b}_i \in \mathcal{E}\}; \mathbf{s} \leftarrow [\mathbf{s}, \mathbf{b}_{:p+1}]; \mathbf{return} \mathbf{s} \rightarrow \mathbf{Early} \mathbf{termination} \mathbf{at} \mathbf{first} \mathbf{end} \mathbf{token}
               31:
848
               32:
                            end if
849
               33:
                            \mathbf{s} \leftarrow [\mathbf{s}, \mathbf{b}]
                                                                                                                                                      ▶ Append full block
850
               34: end while
851
               35: return s
852
```

853 854

855 856

857

858

859 860

862

863

A.5 ILLUSTRATION OF TRAINING LOSS AND ATTENTION DESIGN

Figure 3(a) shows the training loss of our framework: AR loss is applied to text spans, while NAR loss—based on discrete diffusion—is used for audio spans. Although we employ discrete diffusion, our framework is extensible to other NAR generation methods. Figure 3(b) visualizes the attention pattern described in Section 3.4.

A.6 DATASET DETAILS

Table 2 provides a summary of the training datasets, with detailed examples provided in the Appendix A.8.1. During training, we aim to construct a balanced corpus that supports effective learning

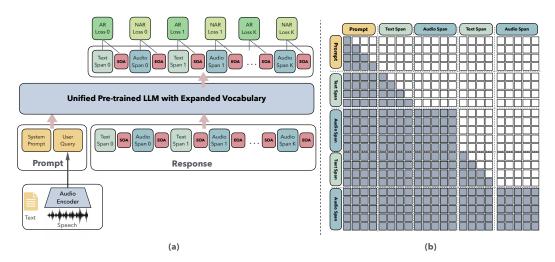


Figure 3: **Training loss and attention design.** (a) **Training pipeline.** Starting from a pretrained text LLM, we expand the vocabulary with audio tokens and control symbols. Text spans use AR cross-entropy loss while audio spans use NAR diffusion loss, sharing a single Transformer backbone. (b) **Attention pattern.** Text spans follow causal attention (left-to-right), while audio spans use bidirectional attention within spans but causal attention across spans, enabling parallel audio generation while preserving cross-modal dependencies.

across multiple tasks. Specifically, we randomly sample one million instances from the ASR dataset, the TTS dataset, and the audio chat dataset respectively. In addition, we create bilingual interleaved text and audio data, ensuring that Chinese and English are represented in approximately equal proportions. To build the audio chat corpus, we rely on the text-to-audio dataset VoiceAssistant-400K together with the text-based datasets OpenHermes-2.5 and Firefly-Train-1.1M, and we employ a TTS model, namely CosyVoice2, to convert text into synthetic audio so as to enrich the training data. To further enhance cross-modal alignment between text and audio, we follow prior work (Zeng et al., 2025) and supplement the training corpus with interleaved text and audio data derived from the large-scale pretrained corpus FineWeb-Edu. This strategy not only expands task coverage but also strengthens the model's ability to jointly learn from and align textual and acoustic modalities. The evaluation datasets are shown in Table 4.

A.7 EVALUATION DEATILS

A.7.1 EVALUATION TASKS

Audio-QA Task Previous evaluation frameworks such as Kimi-Audio assess Audio-QA performance using the text portion of interleaved outputs, which overlooks the fact that the audio output of an end-to-end speech model more directly reflects its ability to generate natural and semantically faithful responses. To address this limitation, we evaluate Audio-QA directly on the audio outputs of our framework by first applying an ASR model to transcribe the generated audio into text, where Whisper-Large-v3 is used for English audio and Paraformer-zh for Chinese audio, with a comparison of ASR performance across different models provided in Table 3. The transcribed text is then combined with the original QA queries and the ground truth answers and passed to a large scale reasoning model, Qwen3-235B-A30B, which serves as an LLM-as-a-Judge model to determine whether the response semantically matches the reference and to provide either a correctness label or a graded score. We report the average accuracy or score on the benchmark, and this evaluation pipeline provides a more faithful assessment of our model's audio-to-audio QA ability in realistic conversational scenarios where speech serves as the output modality.

ASR Task To assess the model's capability in aligning speech with textual representations, we evaluate it on the ASR task, where the model generates text transcriptions from input audio and performance is measured using word error rate (WER). A lower WER indicates more accurate recogni-

Table 2: Summary of datasets used in training.

Dataset	Language	Samples	Task Type	
Emilia_zh	Chinese	500000	TTS	
Emilia_en	English	500000	TTS	
AISHELL2	Chinese		ASR	
AISHELL3	Chinese		ASR	
CommonVoice	Chinese, English		ASR	
GigaSpeech	English	600000	ASR	
LibriSpeech	English	000000	ASR	
MLS-Eng	English		ASR	
PeopleSpeech	English		ASR	
VoxPopuli	English		ASR	
WenetSpeech	Chinese	400000	ASR	
VoiceAssistant-400K	English		Audio Chat	
OpenHermes-2.5	English	1000000	Audio Chat	
Firefly-Train-1.1M	Chinese		Audio Chat	
MathInstruct	English	262039	Text Chat	
MACS	English		AAC	
Clotho-v2	English		AAC	
Nonspeech7k	English	59282	SEC	
VocalSound	English		SEC	
CochlScene	English		ASC	
Chinese-Fineweb-Edu (Skypile)	Chinese	1500000	Interleaved Data	
FineWeb-Edu	English	1500000	Interleaved Data	
Total	_	6321321		

Table 3: WER performance of different ASR models on Chinese (zh) and English (en).

Model	WER-zh (↓)	WER-en (↓)
Whisper-Large-v3	0.5054	0.2167
Paraformer-zh	0.1028	0.3946

tion, which reflects not only strong ASR ability but also effective cross modal consistency achieved by our hybrid AR-NAR modeling framework.

A.7.2 EVALUATION DATASETS

We evaluate model performance on a diverse set of benchmarks covering both Audio Question Answering (Audio-QA) and Automatic Speech Recognition (ASR) tasks. For Audio-QA, we use four datasets: AlpacaEval, TriviaQA, and WebQuestions (English), along with LLaMAQuestions (English), assessing cross-lingual reasoning and comprehension from speech. For ASR, we include five datasets: Fleurs-zh/en (multilingual), AISHELL-1/2, and WenetSpeech (all Chinese), covering varied domains, accents, and recording conditions to robustly measure transcription accuracy. Dataset details are summarized in Table 4.

A.8 TRAINING DETAILS

We train our model using the AdamW optimizer with a global batch size of 2048, a learning rate of $2e^{-5}$, and a weight decay factor of $1e^{-2}$. The learning rate follows a cosine decay schedule with a linear warmup ratio of 0.01. Training incorporates three stochastic strategies: (1) batchwise AR & NAR objective mixing with probability 0.3; (2) prefix preservation masking with ratio 0.3; (3) stochastic span truncation with probability 0.5. During inference, the model alternates between AR text decoding and NAR diffusion-based audio generation, where text decoding uses nucleus sampling with k=10 and p=0.95, and audio spans are generated with 200 diffusion steps, a block

Table 4: Evaluation datasets used	l for Audio-OA and ASR tasks.
-----------------------------------	-------------------------------

Dataset	Language	Task Type	Abbreviation
AlpacaEval	English	Audio-QA	AE.
LLaMAQuestions	English	Audio-QA	LQ.
TriviaQA	English	Audio-QA	TQA.
WebQuestions	English	Audio-QA	WQ.
Fleurs-zh	Chinese	ASR	Fzh.
AISHELL-2	Chinese	ASR	A2.
AISHELL-1	Chinese	ASR	A1.
WenetSpeech-test_meeting	Chinese	ASR	WS_met.
WenetSpeech-test_net	Chinese	ASR	WS_net.
Fleurs-en	English	ASR	Fen.

length of 32 tokens, and a total diffusion span length of 640 tokens under classifier-free guidance with scale 0.1. Since different training strategies lead to varying convergence speeds, reported results are based on checkpoints where training loss has converged. All experiments are conducted on 4 nodes with 8 NVIDIA A100 GPUs per node using the DeepSpeed runtime.

A.8.1 Data Format of Training Data

To enable unified training across diverse tasks, we transform all datasets into a consistent in-put-output format. On the one hand, this standardization allows the model to seamlessly integrate heterogeneous modalities such as speech, text, and interleaved audio-text sequences. On the other hand, a unified design is essential for supporting our training strategies, including batchwise AR & NAR objective mixing, prefix preservation masking, and stochastic span truncation. These strategies rely on a shared representation to operate across modalities in a consistent way. For clarity, we provide representative examples of the adopted data formats as follow, covering ASR, TTS, audio chat, text chat, AAC, SEC, ASC, and interleaved text-audio data.

Figure 4: Example of ASR data format.

Text-to-speech (TTS) Data Format

Figure 5: Example of TTS data format.

Audio Chat Data Format

Figure 6: Example of audio chat data format.

Figure 7: Example of text chat data format.

AAC/SEC/ASC Data Format

Figure 8: Example of AAC/SEC/ASC data format.

Figure 9: Example of interleaved data format.

A.9 USAGE OF LLM

In this paper, the LLM is employed solely for text refinement—correcting typos, fixing spelling errors, and enhancing readability. It is not used for generating research ideas, producing results, or creating original content.